

Testing MapReduce-Based Systems

João Eugenio Marynowski¹, Michel Albonico¹, Eduardo Cunha de Almeida¹, Gerson Sunyé²

¹ Federal University of Paraná, Brazil
`{jeugenio,albonico,eduardo}@inf.ufpr.br`
² INRIA - University of Nantes, France
`gereson.sunye@univ-nantes.fr`

Abstract. MapReduce (MR) is the most popular solution to build applications for large-scale data processing. These applications are often deployed on large clusters of commodity machines, where failures happen constantly due to bugs, hardware problems, and outages. Testing MR-based systems is hard, since it is needed a great effort of test harness to execute distributed test cases upon failures. In this paper, we present a novel testing solution to tackle this issue called HadoopTest. This solution is based on a scalable harness approach, where distributed tester components are hung around each map and reduce worker (i.e., node). Testers are allowed to stimulate each worker to inject failures on them, monitor their behavior, and validate testing results. HadoopTest was used to test two applications bundled into Hadoop, the Apache open source MapReduce implementation. Our initial implementation demonstrates promising results, with HadoopTest coordinating test cases across distributed MapReduce workers, and finding bugs.

Categories and Subject Descriptors: H.3.4 [Information Storage and Retrieval]: Systems and Software

Keywords: MapReduce, Hadoop, Testing, Large-scale

1. INTRODUCTION

Over the last decade the amount of data generated by several applications, like customer reports, survey data and social networks, have reached several petabytes and this amount of data tends to increase dramatically along the next few years up to iotabytes [Winter and Kostamaa 2010]. The analysis of this large amount of data requires a great effort of data processing, drawing a widespread attention from both, the academic community and the industry. Thus, new types of frameworks for large-scale data processing have been proposed. One of them, MapReduce (MR) [Dean and Ghemawat 2004] appears as the most popular one. It allows non-expert users to easily use a large number of machines to process a large amount of data. Among several implementations of MapReduce, the open-source framework Hadoop [The Apache Software Foundation 2011a] from the Apache Foundation has been used by a growing number of companies, such as: EBay, LinkedIn, Quantcast, Facebook and Yahoo! [The Apache Software Foundation 2011b]. Moreover, some DBMS vendors, including Aster, Greenplum and Vertica, have started to integrate MR front-ends into their systems to run on large-cluster machines.

As any other large-scale environment, MR implementations constantly face failures due to several different conditions (e.g., outages, hardware problems, bugs) [Liu et al. 2004]. Thus, these implementations must be designed to be fault-tolerant and tested intensively to ensure they are reliable. However, testing large-scale systems is a complex task and requires a test harness, i.e. a framework to build tests and control their execution by monitoring the behavior of a system under test (SUT) and its outputs [Walter et al. 1998]. The main issue of a test harness is to control the execution of distributed test cases across a large-scale environment. More precisely, this control requires: the

Work partially funded by the Datluge CNPq-INRIA project.

coordination of distributed test cases, the monitoring of the nodes that execute the test cases, and the injection of faults to reproduce real-world conditions. However, available solutions for testing MR implementations are limited and none of them tackle completely this issue.

In this work, we present HadoopTest, a test harness for MR-based implementations. HadoopTest deploys and manages the execution of test cases across distributed nodes. This is achieved by controlling each MR worker (i.e. node) with a distributed *Tester*. A tester is a component that controls the overall test case execution through a distributed coordination algorithm. HadoopTest was used to test two applications bundled with Hadoop. The initial implementation demonstrates promising results, with HadoopTest coordinating test cases across distributed MR workers and finding bugs.

The rest of the paper is organized as follows. The next session introduces the basic concepts of MapReduce and software testing. Section 3 presents our framework for testing MR systems. Section 4 describes the initial results through implementation and experimentation. Section 5 discusses related work. Section 6 concludes.

2. BASIC CONCEPTS

MapReduce has a simplified programming model, leaving aside the treatment of events related to the distributed environment, without programmer intervention. Its programming environment is based on two higher-order functions: Map and Reduce. Each one of these functions has a well-defined behavior. They work together, where the mapping is the initial analysis of the data input and the output is handled by the reduce function. Both the map and the reduce functions are programmed by the user. Several copies of the program are distributed across machine-nodes, where one copy is called master and the other copies are workers. The master selects idle workers to assign a map or a reduce instance. The data flow between the map and the reduce functions is shown in Figure 1.




Fig. 1: Execution of Map and Reduce operations

All processing is done based on $\langle key, value \rangle$ pairs. For each key, the map function associates a value. The user defines the correspondence between the *value* and *key*. The output of the map function is said to be the intermediate key pair.

For each range of keys will be assigned an intermediate *Reducer*, a service responsible for performing the reduce function. As a result, the operation will generate a pair $\langle key, result \rangle$. Analogous to the map function, it is up to the user to define the correspondence between the *key* and result. When the map function terminates, the reduce function starts. When all the reduce instances terminate, they append their result to a final output file.

2.1 Testing Background

Testing is a process to validate a system to ensure its quality [Ammann and Offut 2008]. In our experiments, we focus on functional testing, whose principle is to apply inputs to the SUT and to compare the observed outputs to the expected results. This comparison is made by an *oracle*, the mechanism responsible for assigning a *verdict*. If the expected results and the observed outputs are the same, the verdict is *pass*. Otherwise, the verdict is *fail*. The verdict may also be *inconclusive*, meaning that the test result is not precise enough to satisfy the test criteria.

There are different sorts of oracles [Baresi and Young 2001], for instance, assertions, value comparison, log file analysis, and manual analysis. Our approach is based on assertions and value comparison due to their simplicity. Both types of test oracles are sufficient to the HadoopTest initial implementation, since we are testing simple MR applications. However, we plan to add a broader range of test oracles in the future.

3. TESTING MAPREDUCE-BASED SYSTEMS

Testing MR-based systems is a complex task, for which the test harness must execute distributed test cases. The main issues of the testing harness are: the coordination of distributed test cases, the control of distributed workers and the injection of faults.

Distributed testing is only possible through a mechanism to coordinate the execution of distributed test cases across distributed nodes. In the context of MR, this can be more complex, since several workers with different functionalities are running together. For instance, while a worker is applying a map function another is leaving the system due to a fault. This also leads to an individual control of the workers, which is important to put the workers in any state along testing.

The fault injection issue is important to reproduce a real-world environment during testing. Several types of failures must be considered, such as: node failure, network traffic, high latency and performance variations. In addition, these faults may be executed individually at each worker. For instance, to reject the disk access to a specific reduce instance to test fault-tolerance.

3.1 HadoopTest

In this section, we present HadoopTest, a test harness solution for MR-based systems. This solution is based on the individual control of distributed MR workers across large-cluster machines. This approach allows to coordinate the execution of test case actions on different workers in parallel. HadoopTest allows the combination of fault injection and functional tests to build complex test cases. Test cases deployed on HadoopTest are not intrusive, meaning that API calls are not added to the SUT source code. The reason is that additional code may add bugs to the SUT, making it difficult to detect the source of a failure, which could be either the SUT code or the inserted code.

Test cases are written in Java, similarly to JUnit [jun 2011], to accelerate HadoopTest acceptance, although we plan to improve the testing language in the future (e.g., using Bloom [Alvaro et al. 2010]). More precisely, test cases are Java classes and test case actions are instance methods, marked with an annotation. Annotations are metadata where test engineers define deployment instructions and the test case execution workflow. HadoopTest uses reflection to read this metadata to achieve coordinating the test cases (detailed in Section 3.1.1).

The HadoopTest architecture, shown in Figure 2a, is composed of a test *coordinator* and distributed *testers*. The coordinator is responsible for three main tasks: (1) to control the execution of distributed test cases, (2) to coordinate the test case actions and (3) to gather the verdicts from the testers. The testers receive coordination messages from the *coordinator* to execute the test case actions on the master and distributed workers. These messages can also contain stimulus to inject faults on the

workers. In the present version of HadoopTest, we only implemented outages to drop workers. Other types of failures depend on a lower-level control of the SUT, for instance, disk or network failures. Once the testers reach the end of a test case execution, they validate the results to assign a local verdict. This validation is based on assertions and value comparison.

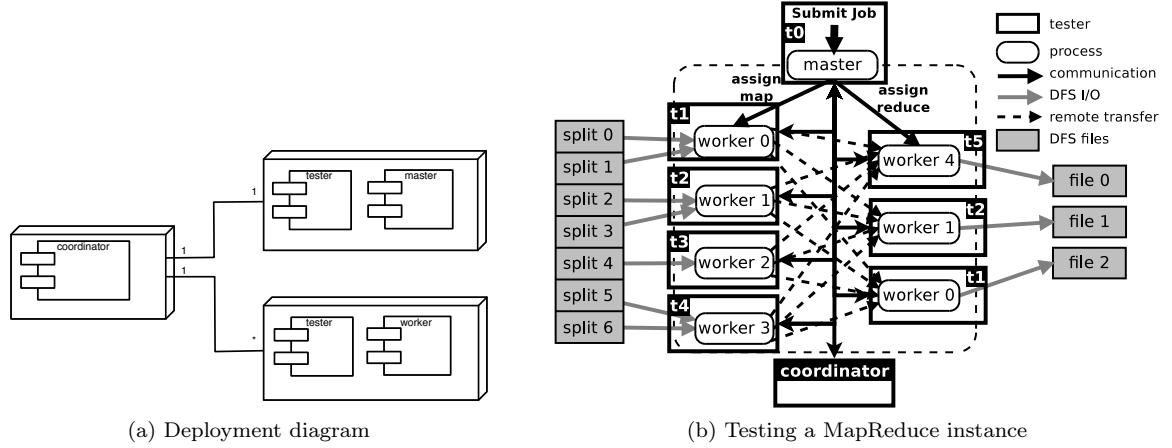


Fig. 2: HadoopTest diagrams

Figure 2b shows the application of HadoopTest to an MR instance. The *coordinator* controls the execution of six testers, identified by $t_0..t_5$. Tester t_0 controls the master and each one of the other testers, $t_1..t_5$, controls a worker instance.

3.1.1 The Coordination of Actions. We let $\tau = (A^\tau, V^\tau, T^\tau)$ denote a distributed test case, where A^τ is a set of test actions, V^τ is a set of local verdicts, and T^τ is a set of testers. A test case action is a tuple $a = (T', I, \theta, h)$ where T' is a set of testers that should execute the action, I is a set of instructions (e.g., Java instructions), and θ is the interval of time in which the action should be executed. This interval is important to bound the execution time, since coordination messages are asynchronous. We let h denote the hierarchical order that the actions will execute. Naturally, actions with the same hierarchical order are allowed to execute in parallel.

The overall coordination approach is outlined in Algorithm 1. Initially, the coordinator registers into a schedule S^τ each tester t and the test actions A^τ that such tester will invoke. The schedule S^τ maps actions onto sets of testers. Once the registration is finished, one coordination message is sent to each corresponding tester following the sequence of actions. The coordination messages are sent asynchronously to allow the execution of the actions in parallel. Local verdicts are also received asynchronously from the testers. If an action takes more than θ , then the verdict is *inconclusive*. At the end, the oracle returns the final verdict that gathers all the local verdicts.

3.1.2 Test Case Example. Table I shows a simple distributed test case to illustrate the algorithm. The goal is to validate a computation upon failures. This test case involves six testers $T = \{t_0, \dots, t_5\}$ and seven test actions $A = \{a_0, \dots, a_6\}$. The tester t_0 executes the action a_0 to start the master. The testers $\{t_1, \dots, t_5\}$ execute the action a_1 to start the workers. Next, tester t_0 submits a job to the master at action a_2 . This job is composed by a computation in the form of map and reduce functions. During the execution of the job, the tester t_2 is dropped from the system. The validation is done at action a_4 . If the output data is the same as the expected, then the verdict is *pass*. Otherwise, the verdict is *fail*. If t_0 is not able to retrieve any output data, then the verdict is *inconclusive*. The rest of the actions are executed to stop the SUT.

Algorithm 1: Coordination Algorithm

Input: A^τ , an ordered set of test actions; T^τ , a set of testers;
Output: a global verdict
 $S^\tau \leftarrow \emptyset$;
 $V^\tau \leftarrow \emptyset$;
forall the $t \in T^\tau$ **do**
 $S^\tau \leftarrow \text{register}(t, A^\tau)$;
forall the $a \in A^\tau$ *such that h is the same* **do**
 Send coordination messages for all $t \in T'$ such that $a \mapsto T'$ up to θ ;
 if θ **then**
 $V^\tau \leftarrow \text{inconclusive}$;
 else
 $V^\tau \leftarrow t.\text{result}(a^h)$;
return $\text{oracle}(V^\tau)$;

Table I: MapReduce simple test case

Action	Hierarchy	Tester	Instructions	Timeout
a_0	0	t_0	startMaster()	100
a_1	1	t_1-t_5	startWorkers()	1000
a_2	2	t_0	sendJob()	1000000
a_3	2	t_2	dropWorker()	1000
a_4	3	t_0	assert(output, expected)	10000
a_5	4	t_1, t_3-t_5	stopWorkers()	1000
a_6	5	t_0	stopMaster()	1000

4. EXPERIMENTAL VALIDATION

This section presents the evaluation of HadoopTest through experimentation. First, we evaluated the overhead produced by HadoopTest to coordinate the execution of distributed test cases. Second, we used HadoopTest to test two applications bundled into Hadoop, however, due to the lack of space, we only present the PiEstimator results. The objective is to validate whether HadoopTest is able to identify bugs.

The PiEstimator calculates the π value using the Monte Carlo method, that considers a circle exactly inscribed inside a square with side length 1. The map function randomly creates points inside the square and counts the points placed inside and outside the circle. The reduce function accumulates the points inside (I) and outside (O) counted by the map function. The estimated π value is obtained by $4 * (I/T)$, where $T = I + O$.

The experiments were executed on the Grid5000 platform¹ using up to 50 cluster machines running Debian GNU/Linux. The cluster machines were connected by a 1 Gbps network and they had a similar configuration: 2 Intel Xeon 2.6GHz dual-core processors, 8 GB RAM memory and 250 GB SATA HD.

4.1 The HadoopTest Overhead

To evaluate the overhead produced by HadoopTest, we executed the PiEstimator application in two ways. In the first one, Hadoop is executed alone, in order to evaluate the raw execution time. In the

¹Grid 5000 Platform: <http://www.grid5000.fr>

second one, Hadoop is executed along with HadoopTest, in order to evaluate the overhead produced during testing. We perform tests with 2, 10, 20, and 50 machine-nodes on the Grid5000. Figure 3 shows the execution time of PiEstimator running on Hadoop and HadoopTest with 50 machine-nodes. We vary the number of map instances in each execution.




Fig. 3: Execution time variance of the PiEstimator

The execution with HadoopTest presented an average overhead of 30%. Although this overhead decreases the performance of the tests, it did not harm the test results. In fact, this overhead is important while testing applications where timing constraints are important. Future work is needed to reduce the coordination messages and consequently, reducing the overhead. Nevertheless, HadoopTest coordinated up to 10,000 map instances concurrently, however, the overhead due to the number of threads running along was overwhelming to 50 machine-nodes. A larger number of machine-nodes is required to test this number of map instances.

4.2 The HadoopTest Results

To evaluate whether HadoopTest is able to identify bugs, we used Mutation Analysis [Ma et al. 2006; Offutt 1994] to create a set of buggy versions (i.e., mutants) of the PiEstimator. We used the ASM [Bruneton et al. 2002] Java bytecode engineering library to create mutants. In our experiments, a mutation is a modification of an arithmetic and/or a logic operator into the original bytecode to generate an incorrect result. The goal is to identify the largest possible number of incorrect results.

We generated 13 mutants of the PiEstimator class and the results are shown in Table II. The expected π value returned by the original application was 3.1416 and only the mutants M1, M6, M7, M9 and M12 returned this value. These mutants have the *Pass* verdict on the test case execution while the other mutants M4, M5, M10 and M11 received a *Fail* verdict, since the π computation parameters were modified resulting in a different value than the expected one. In the case of the mutants M0, M2, M3 and M8, the modifications were in the execution parameters which interfered on their correct execution. They returned NULL as results.

HadoopTest effectiveness was evaluated in terms of the number of detected mutants. When mutation analysis is applied to a program code and generates several mutants, some of them are equivalents to the original source for different reasons: the modified part is never executed, binary operators

Table II: Test case over 13 mutants from the PiEstimator

Mutants	Result	Pass	Fail
M0	NULL		X
M1	3.1416	X	
M2	NULL		X
M3	NULL		X
M4	3.0776		X
M5	3.1312		X
M6	3.1416	X	
M7	3.1416	X	
M8	NULL		X
M9	3.1416	X	
M10	3.1408		X
M11	3.1408		X
M12	3.1416	X	

are equivalents for this precise case, etc. Here, we considered that the mutants that obtained the same output as the original application are equivalents. The initial implementation of HadoopTest demonstrated promising results by identifying all of the non-equivalent mutants of the PiEstimator. This was also true for the WordCount application tests.

5. RELATED WORK

MRUnit [mru 2011] and Herriot [her 2011] are the most popular tools to test MR implementations. They are designed to help bridging MR implementations and JUnit testing tool. Both tools provide an API that must be used along the development of an MR-based system. Differently from MRUnit, Herriot provides an external component, Test Node, which is responsible for the overall test execution control. This control is only possible including some API calls into the SUT source-code. However, this intrusive approach may spoil the source code generating new bugs. Differently from our approach, both tools neither coordinate different nodes at the same time nor stimulate them in parallel which restricts to build complex test cases.

Ganesha [Pan et al. 2009] is a black-box testing tool used to detect performance problems in MR systems. Differently from MRUnit and Herriot, it does not require any modification on the SUT source code. Ganesha differs from our approach in the validation aspect. It monitors both, the OS-level and Network-level counters to detect the culprit node that compromised the MR performance. This approach produces low performance overhead, however, the validation is narrowed by the limitations of the OS-level and Network-level counters. In addition, Ganesha requires an additional effort to filter the metrics of concurrent applications out of the counters (including MR's). In contrast, HadoopTest provides on-line assertion mechanisms that can be used by any kind of test (including performance test) as well as the individual control of nodes, which can be used to detect any performance problem.

PeerUnit [de Almeida et al. 2010; de Almeida et al. 2010] is a framework for testing peer-to-peer systems. It ensures the sequencing of test case actions synchronizing their execution through distributed nodes. PeerUnit allows to coordinate any SUT component also leveraging the approach of distributed testers. It differs from our approach on the test case execution workflow. First, PeerUnit can not execute different actions in parallel, preventing the writing of complex test cases. Second, nodes (i.e., peers) are considered functionally identical, which is not the case of MR implementations, where they are divided into: workers and master nodes (as detailed in Section 2). Moreover, different workers may perform different actions along the execution of an MR application. In our approach, distributed testers are designed to achieve handling both nodes and actions in parallel.

6. CONCLUSION

In this paper we presented HadoopTest, a test harness solution for testing MapReduce-based systems. HadoopTest tackles two important issues of MapReduce testing. First, it coordinates the execution of test cases across distributed MapReduce workers. This is achieved by controlling each MR worker with distributed testers. Second, HadoopTest allows the coordination of the parallel execution of test case actions on different MR workers. Naturally, in practice it allows to combine fault injection with functional tests to built complex test cases.

Through experimentation on 50 machine-nodes on the Grid5000, we showed that HadoopTest coordinates distributed test cases across MR workers. We also used HadoopTest to test two applications bundled into Hadoop, the PiEstimator and the WordCount. We used mutation analysis to create mutants of these applications. HadoopTest demonstrated promising results by identifying all of the non-equivalent mutants.

In the future, we plan to execute more complex experiments involving different fault injections. This can be achieved by instrumenting the underlying structure of the SUT (e.g., network, disk), which is not yet possible with this version of HadoopTest. In addition, we plan to test more complex MapReduce applications, such as: Hive, Mahout, and Nutch. These applications provide a broader range of functionalities to test.

REFERENCES

- Herriot, large-scale automated test framework. <https://issues.apache.org/jira/browse/hadoop-6332>, 2011.
- Junit website. <http://www.junit.org>, 2011.
- Mrunit project. <http://archive.cloudera.com/docs/mrunit/index.html>, 2011.
- ALVARO, P., CONDIE, T., CONWAY, N., ELMELEEGY, K., HELLERSTEIN, J. M., AND SEARS, R. Boom analytics: exploring data-centric, declarative programming for the cloud. In *EuroSys*, C. Morin and G. Muller (Eds.). ACM, pp. 223–236, 2010.
- AMMANN, P. AND OFFUT, J. *Introduction to Software Testing*. Cambridge University, 2008.
- BARESI, L. AND YOUNG, M. Test oracles. Technical Report CIS-TR-01-02, University of Oregon, Dept. of Computer and Information Science, Eugene, Oregon, U.S.A. August, 2001. <http://www.cs.uoregon.edu/~michal/pubs/oracles.html>.
- BRUNETON, E., LENGLET, R., AND COUPAYE, T. Asm: A code manipulation tool to implement adaptable systems. In *In Adaptable and extensible component systems*, 2002.
- DE ALMEIDA, E. C., MARYNOWSKI, J. E., SUNYÉ, G., AND VALDURIEZ, P. Peerunit: a framework for testing peer-to-peer systems. In *ASE*, C. Pecheur, J. Andrews, and E. D. Nitto (Eds.). ACM, pp. 169–170, 2010.
- DE ALMEIDA, E. C., SUNYÉ, G., TRAON, Y. L., AND VALDURIEZ, P. Testing peer-to-peer systems. *Empirical Software Engineering* 15 (4): 346–379, 2010.
- DEAN, J. AND GHEMAWAT, S. Mapreduce: Simplified data processing on large clusters. In *OSDI*. pp. 137–150, 2004.
- LIU, Y., LIU, X., XIAO, L., NI, L. M., AND ZHANG, X. Location-aware topology matching in p2p systems. In *INFOCOM*, 2004.
- MA, Y.-S., OFFUTT, J., AND KWON, Y.-R. Mujava: a mutation system for java. In *Proceedings of the 28th international conference on Software engineering*. ICSE '06. ACM, New York, NY, USA, pp. 827–830, 2006.
- OFFUTT, A. J. A practical system for mutation testing: Help for the common programmer. In *ITC*. IEEE Computer Society, pp. 824–830, 1994.
- PAN, X., TAN, J., KAVULYA, S., GANDHI, R., AND NARASIMHAN, P. Ganesha: blackbox diagnosis of mapreduce systems. *SIGMETRICS Performance Evaluation Review* 37 (3): 8–13, 2009.
- THE APACHE SOFTWARE FOUNDATION. Apache hadoop. <http://hadoop.apache.org/>, 2011a.
- THE APACHE SOFTWARE FOUNDATION. Apache hadoop users. <http://wiki.apache.org/hadoop/PoweredBy/>, 2011b.
- WALTER, T., SCHIEFERDECKER, I., AND GRABOWSKI, J. Test architectures for distributed systems: State of the art and beyond. In *IWTCS*, A. Petrenko and N. Yevtushenko (Eds.). IFIP Conference Proceedings, vol. 131. Kluwer, pp. 149–174, 1998.
- WINTER, R. AND KOSTAMAA, P. Large scale data warehousing: Trends and observations. In *ICDE*, F. Li, M. M. Moro, S. Ghandeharizadeh, J. R. Haritsa, G. Weikum, M. J. Carey, F. Casati, E. Y. Chang, I. Manolescu, S. Mehrotra, U. Dayal, and V. J. Tsotras (Eds.). IEEE, pp. 1, 2010.